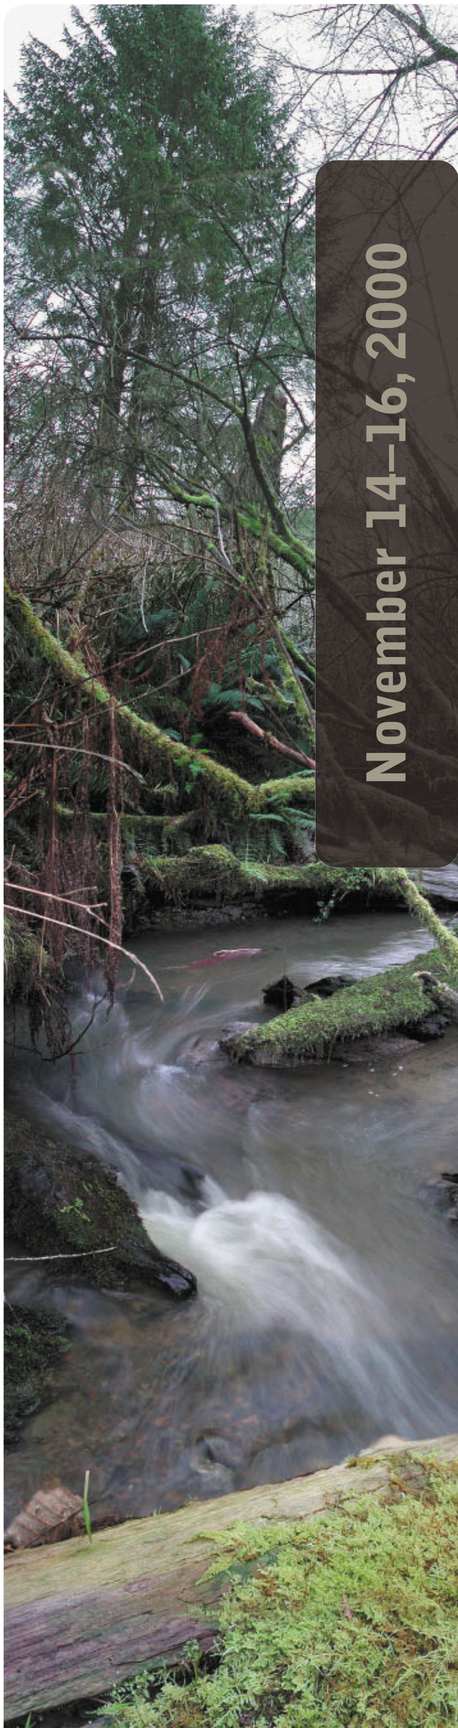

Proceedings of the Salmon Habitat Restoration Cost Workshop

November 14–16, 2000

Oxford Suites Hotel
Gladstone, Oregon

Proceedings of the Salmon Habitat Restoration Cost Workshop

Stan T. Allen, Editor

Cindy Thomson, Co-Editor

Robin Carlson, Co-Editor

Published by

Pacific States Marine Fisheries Commission

205 SE Spokane Street, Suite 100

Portland, Oregon 97202-6413

Tel: (503) 595-3100 Fax: (503) 595-3232

March 2004

graphic design: Jeff Bright, www.jeffbrightdesign.com

spawning coho salmon photos: Thomas Dunklin, www.thomasbdunklin.com

Conceptual framework(s) for cost analysis

The Allocation Problem in Habitat Restoration

DAVID TOMBERLIN
National Marine Fisheries Service
Southwest Fisheries Science Center
110 Shaffer Road
Santa Cruz, CA 95060
David.Tomberlin@noaa.gov

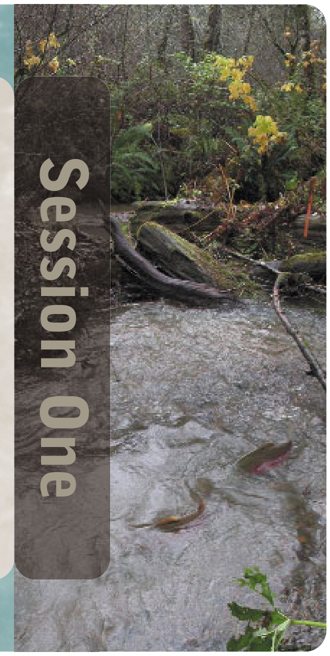

ABSTRACT

This paper explores the question of how best to allocate habitat restoration effort over space and time. Stylized examples are used to illustrate how threshold effects, competition among projects, risk, learning, and the choice of restoration objective affect desirable effort allocations. The paper concludes with some thoughts on the applicability of decision modeling to the habitat restoration planning problem.

INTRODUCTION

The allocation of limited resources over space and time is a key element of habitat restoration planning. Restoration planners must choose which activities to undertake, whether to spread effort among many projects or to focus on a few projects, and whether to launch projects as quickly as possible or to proceed experimentally. This paper identifies some salient features of this allocation problem, demonstrates their influence on desirable effort allocations, and assesses the suitability of decision modeling techniques for restoration planning.

The goal of habitat restoration may be expressed in general terms, such as "recover endangered species" or "improve habitat," but here the emphasis will be on goals that can be expressed as optimization problems, such as "maximize the number of returning spawners" or "minimize extinction risk for a population." The goal may be expressed in terms of restoration activity (e.g., miles of road decommissioned), human values (e.g., social welfare or economic impact), fish population characteristics (e.g., population size or extinction risk), or ecosystem characteristics (e.g., temperature change or reduced sediment load). Each of these can imply a different best allocation of restoration effort. This paper treats ecosystem characteristics as the yardstick by which success is measured and the level of restoration activity as the choice variable. Of course, it may in practice be very difficult to assess how ecosystem characteristics change in response to restoration effort.

Other aspects of the decision problem may be as important to the allocation decision as the chosen goal. Allocation decisions must be made at several spatial scales, and an allocation may be efficient in each of its parts yet inefficient as a whole. Cumulative (or threshold) effects within watersheds are likely, so that benefit is not a simple linear function of effort. Temporally, the possibility of learning from pilot projects must be weighed against the potential costs of waiting, and there are often lags between project implementation and efficacy. Metapopulation dynamics are both a spatial and a temporal complication. Importantly, decisions must be made under imperfect information about — or even ignorance of — both natural and social aspects of the project.

Below, stylized examples illustrate the importance of these considerations to the allocation problem. The next section addresses the spatial allocation problem in simple terms through a model with two competing projects. The paper then explores risk due to the inherent uncertainty of project outcomes, the possibility of learning from pilot projects, and the impact of an extremely risk-averse objective function on the desired effort allocation. The concluding section assesses the applicability of decision modeling to habitat restoration planning and suggests some elements of a decision science research to support this planning.

TWO RIVERS, ONE BUDGET[1]

Consider the problem of allocating restoration effort among two identical river basins so as to maximize some measure of ecosystem health. In each river, the more effort expended the greater is ecosystem health, but let us suppose that this relationship is sigmoidal rather than linear — that is, the marginal benefit of restoration effort is small at low effort levels, increases rapidly over the mid-range of effort, and decreases again as ecosystem health tapers off to some ceiling (Figure 1).

Suppose the goal is to maximize the sum of ecosystem health in the two basins, and that effort can be distributed between the

Figure 1. Cumulative effect

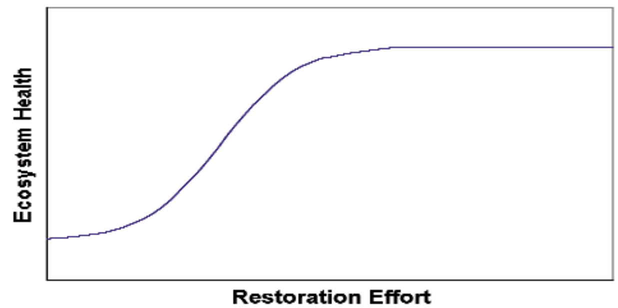

rivers in any way, subject to a limit on total effort (i.e., a 'budget'). Letting H_i and E_i represent health and effort, respectively, in river *i*, the problem can be expressed as one of constrained optimization:

maximize $H_A + H_B$ (1)

subject to 1) $E_A + E_B \leq$ Budget (of time, personnel, funds, etc.)
2) H_i is the sigmoidal function of E_i shown in Fig. 1

Figures 2 and 3 depict this problem graphically by showing the health of the two rivers as mirror images — that is, E_A increases from the left and E_B from the right — which allows the sum $H_A + H_B$ to be depicted at all allocations of effort possible under a given budget constraint (here 12 and 6 units of effort, respectively). Figure 2 shows that when the budget is high relative to what's required to achieve maximum health, it is best to split the effort evenly among the rivers. By contrast, Figure 3 shows that when the budget is relatively small, it is best to concentrate the effort on one river. Given this problem, then, if the budget is very large, many distributions will be optimal or nearly so — but if the budget is small, spreading it among projects is a very poor allocation.

The problem as posed is not realistic, but the simple model provides a clear picture of how rules of thumb (e.g., "address the worst

1- This section follows Wu and Boggess (2000), who analyze the link between fish populations and habitat restoration efforts in the John Day River of Oregon.

Figure 2. Competing projects — large budget case

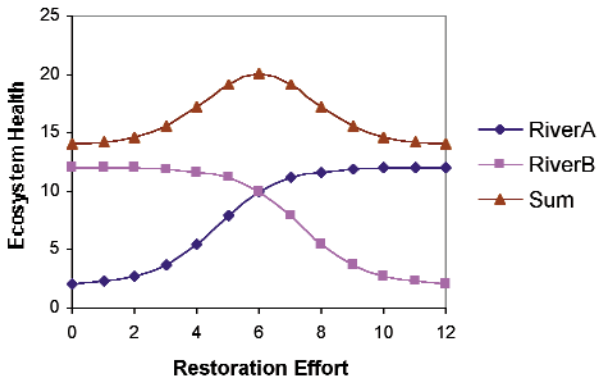

Figure 3. Competing projects — small budget case

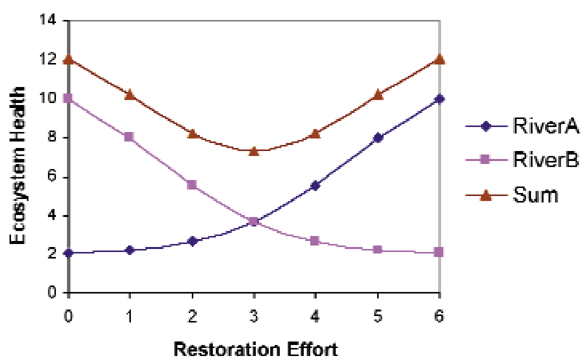

problems first", or "spread the budget widely among projects") may lead to ineffective resource use. More complicated scenarios (heterogeneous rivers, differential responsiveness to restoration effort, etc.) can be formulated as linear or nonlinear programming problems to suggest allocations and to test the sensitivity of different plans to assumptions about model parameters.

INTRODUCING UNCERTAINTY

A serious weakness of the above analysis is that it ignores uncertainty about project outcomes. If outcomes are uncertain in the small-budget case described above, for example, diversification across projects might be desired to reduce the chance of an entirely disastrous outcome, even though we have seen that when outcomes are known with certainty all effort should be expended on a single river. The degree to which outcome risk influences the preferred course of action depends on both the degree of uncertainty about outcomes and on the decisionmaker's attitude toward risk.

A straightforward way to incorporate risk in the analysis is a mean-variance objective function, commonly used for portfolio analysis (see, for example, Bodie, Markus, and Kane 1996)[2]. In this approach, the decision problem is to maximize the expected value of the outcome less some penalty for variability, which is a function of the variance of the sum of ecosystem health in the rivers and a penalty parameter *k*:

$$\text{Maximize} \quad E(H_A + H_B) - kV(H_A + H_B) \qquad (2)$$

subject to the same constraints as before. Here, the parameter *k* represents the decisionmaker's degree of aversion to risk.

Figure 4 shows one river's ecosystem health as a function of restoration effort expended on an uncertain project for which two discrete outcomes are possible. While the expected value (the dashed line) is the same as the deterministic values in Figure 2, here two outcomes are possible, High or Low. Given this uncertainty in two rivers, the decisionmaker must choose how to allocate effort.

Figure 4. Discrete uncertain outcomes — small budget case

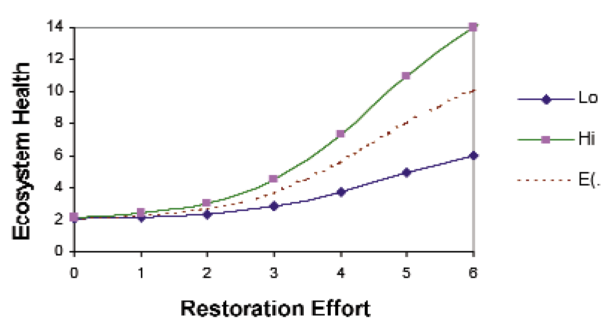

Figure 5 shows the value of the objective (2) under different values of the risk aversion

2- This approach is based on a particular implicit utility function (see e.g., Varian 1993 pp. 189-90).

Figure 5. E(A+B) - kV(A+B) — small budget case

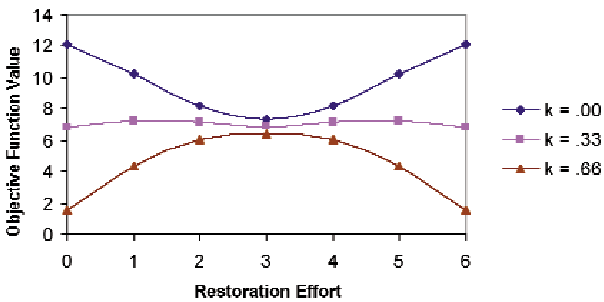

parameter k. The objective value, which is now a function of ecosytem health, increases from the left for River A and from the right for River B, as before. When the decision-maker doesn't care about risk ($k = 0$), the best strategy is to focus all effort on one river, just as in the no-risk case from the previous section; when k is high, it is preferable to spread the funds among projects, that is, to hedge bets.

Quadratic programming, with the objective of minimizing the variance of total benefit subject to some minimum level of benefits, is a convenient way to solve more general problems of this sort. Extinction risk can also be addressed in the allocation problem via an appropriately articulated goal.

LEARNING FROM PILOT PROJECTS

Pilot projects, field trials or other information gathering may reduce uncertainty in many cases, allowing managers to make more informed decisions. The gains from learning must be weighed against the potential cost of waiting, and usually neither can be known with certainty.

The economics literature on learning is extensive (and, by the way, similar to the ecology literature on optimal foraging). Decision trees provide one simple tool for assessing the value of experimentation. Imagine the same uncertainty over outcomes as in the previous section, except that the planner now faces the same allocation decision in two consecutive years, and has the option to reduce uncertainty by learning: by investing at least one unit of effort in a river, the planner can learn the river's "type," which may be either High (i.e., very responsive to restoration work) or Low. Suppose knowing a river's type removes all uncertainty about how it will respond to restoration effort. While knowing the type of both rivers could aid the planner, the information comes at a cost if the planner is risk-neutral, because it can only be obtained by foregoing the option to concentrate effort entirely on one river (recall from Figure 5 that the best decision in the small-budget case, given risk neutrality, was to concentrate on one river). Figure 6 shows a simple example of a decision tree when outcomes are described by the sigmoidal function of previous sections, with uncertain results and a small budget, as in Figure 5. The left-most box represents the time at which the planner chooses whether to invest at least some effort in both rivers so as to learn both their types. The circles represent probability nodes, the planner's estimates of the probabilities of various outcomes depending on the decision made in the first period. In period one, the planner chooses whether to invest in learning, which implies a lower first-period expected benefit but reveals the types of both rivers. After this decision is made, the planner acquires the first-period benefit (10.2 or 12.1) and learns the type of one or both rivers. In period two, if the planner knows the type of both rivers, there will be no uncertainty about the best allocation, and the planner can simply choose the best strategy depending on the information now available. Depending on whether the types are both High, both Low, or mixed, the second-period benefits are 16.2, 8.1, or 16.1, respectively. By contrast, if the planner did not invest in information gathering in period 1, but chose to concentrate all effort on River A, the optimal choice will depend on the revelation

Figure 6. A decision tree. This figure shows, from left to right, actions and probabilities of uncertain outcomes before action is taken in period 1. If the planner invests in learning the types of both rivers, there is no uncertainty left in period two. If the planner does not make this investment and learns the type of only one river (here River A), there is still uncertainty about River B's type, but this node is not shown since it is not needed for the planner to make the optimal decision. The planner's problem is to choose the action plan that maximizes the expected benefit of restoration. As the results at the bottom show, in this case not learning turns out to be a better plan (in expectation).

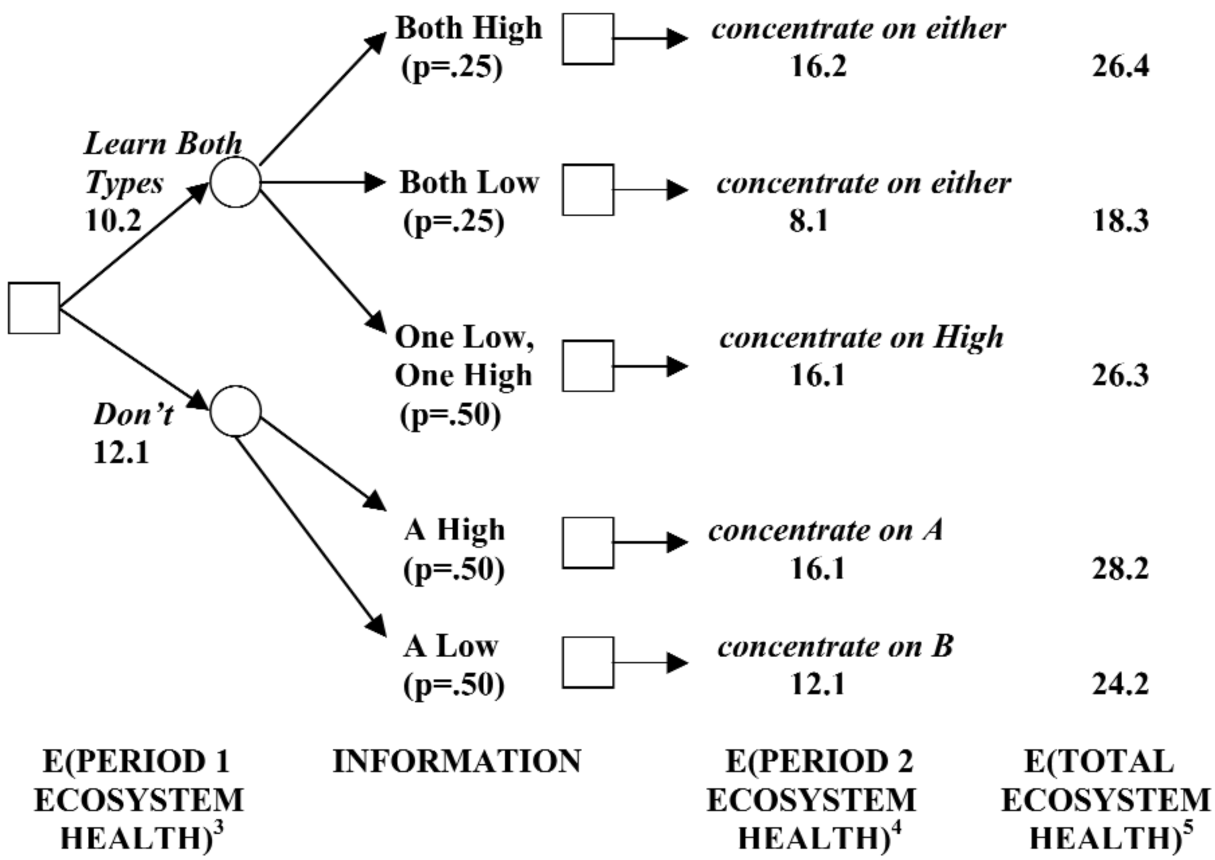

E(Ecosystem Health | Learn)[6] = 24.3

E(Ecosystem Health | Don't Learn)[7] = 26.2

3- Derived from the risk-neutral scenario (k=0.00) in Figure 5: 12.1 corresponds to the highest achievable level of ecosystem health when restoration is concentrated on one river; 10.2 corresponds to the highest achievable level of ecosystem health when at least one unit of the budget is spent on each river, which is necessary to learn the type of both rivers.

4- Derived from Figure 4 for the small budget scenario (total restoration effort = 6) as follows:
16.2 = 14.0 (ecosytem health associated with applying effort=6 to either river of High type)
+ 2.2 (ecosytem health associated with applying effort=0 to the other river of High type)
8.1 = 6.0 (ecosytem health associated with applying effort=6 to either river of Low type)
+ 2.1 (ecosytem health associated with applying effort=0 to the other river of Low type)
16.1 = 14.0 (ecosytem health associated with applying effort=6 to the river of High type)
+ 2.1 (ecosytem health associated with applying effort=0 to the river of Low type)
16.1 = 14.0 (ecosytem health associated with applying effort=6 to River A, type High)
+ 2.1 (expected ecosytem health of applying effort=0 to River B, type unknown)
12.1 = 2.1 (ecosytem health associated with applying effort=0 to River A, type Low)
+ 10.0 (expected ecosytem health of applying effort=6 to River B, type unknown).

5- Sum of Period 1 and Period 2 ecosystem health.

6- 24.3 = 0.25*26.4 + 0.25*18.3 + 0.5*26.3.

7- 26.2 = 0.5*28.2 + 0.5*24.2.

of A's type. If A is revealed to be 'High', the best thing to do is sink all of the next period's budget into A, since it is a sure thing. If A turns out to be 'Low', the best thing to do is bet that B will be 'High'. The decision tree shows the set of possible actions and their expected outcomes. In this (entirely artificial) case, the expected value of learning does not merit the cost, so the planner would be better off (in expectation) to plunge ahead without learning the type of both rivers.

While this example assumes that investing in learning leads to definitive results, the learning process will usually only produce new, but still uncertain, estimates of outcome probabilities. It may even be the case that nothing is learned. Appropriate modification of the decision tree can address these complications. Risk-aversion can also be incorporated into a decision tree via introduction of an appropriate utility function (see Winston 1994 for an accessible introduction). Other analytical techniques to address intertemporal planning problems include dynamic programming (essentially a generalized decision tree representation) and simulation.

DIFFERENT GOALS, DIFFERENT ALLOCATIONS

The scenarios above have assumed that the desired effort allocation maximizes the expected value of ecosystem health, with perhaps a penalty for risk. Another type of goal, often motivated by strong risk aversion or by lack of information, is to maximize ecosystem health in the worst-case scenario (known as the "maximin" strategy). This strategy implies that the decision-maker does not consider the upside potential of risk, but is instead focused entirely on avoiding very bad outcomes.

Figure 7 depicts such a strategy[8]. The choice set is exactly the choice set under uncertainty when both rivers turn out to be type Low (that is, in the worst-case scenario). The goal is to choose an allocation that will

Figure 7. Maximin strategy — small budget

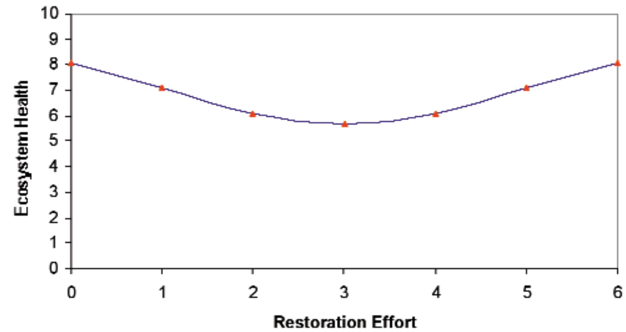

make the best of this worst situation. Given the small budget and this objective, the best we can do is concentrate on one river or the other. Doing so allows us to avoid the worst possible outcome, which would result from splitting effort equally among two rivers that both turn out to be the Low type.

The maximin strategy does not require that outcome probabilities be known or estimated, as long as the set of possible outcomes can be defined. It may be useful as a way for planners to formalize notions of the safe minimum standard (from the economics literature) and the precautionary approach (from the conservation biology literature). In addition to maximin and maximizing expected value of the outcome, many other formulations of the goal can be considered (see Winston 1994 for some examples).

CAN DECISION MODELING CONTRIBUTE TO HABITAT RESTORATION PLANNING?

Restoration planning without reference to a well-posed decision problem may result in significant missed opportunities. For example, failure to account for a nonlinear relation between restoration effort and benefits, such as in the first example above, could lead to greatly reduced restoration efficacy. In the context of habitat restoration planning, a well-posed decision problem should explicitly address threshold effects, budget limitations, risk, and opportunities for exper-

8- Derived by graphing the Low outcome from Figure 4 for Rivers A and B (with HA increasing from the left and HB from the right), then taking their sum.

imentation, if at all possible. A clear statement of the objective is also essential, because different objectives, even when they share the same general aim of conserving endangered species, may result in quite different preferred effort allocations. The modeling approach described in this paper requires that a link between effort and outcomes be established, at least probabilistically. Without this link, there is little basis for taking decisions.

Assuming a coherent decision problem can be developed, what might habitat restoration planners gain by using decision models of the sort described above? In practical terms, models may produce situation-specific results or aid in the development of rules of thumb that can be used across projects, watersheds, and populations. Conceptually, modeling can address uncertainty, scale issues, extinction risk, and the incorporation of information through Bayesian learning, multiple objectives, multiple inputs, and multiple outputs. While models that produce precise prescriptions are almost surely unattainable in the field of habitat restoration planning, modeling can bring focus to planning discussion and help prioritize information gathering needs.

Mathematical programming models drawing on the many tools available in the operations research literature have been stressed above. Analytical models can also provide some useful insight into the nature of the problem, and with a range of plausible parameters can test whether there may be some reasonably robust rules of thumb for allocation. Simulation models may be helpful when the complexity of the system makes mathematical programming intractable.

The stylized examples presented above suggest a research agenda, focusing on information and risk, for the decision science aspect of habitat restoration planning. Key elements of this agenda might include:

• Representing the "technology" of habitat restoration, such as substitutability or complementarity of activities in habitat restoration, which could imply very different efficient effort allocations.

• Models to identify effective information acquisition strategies.

• Linking a restoration effort allocation model to a model of extinction risk, which could enable planners to address, for example, the feasibility of mitigating habitat loss by increasing restoration effort on other lands.

• Providing a framework for considering trade-offs among risks (short-term vs. long-term risks, risks in one population vs. in another).

• Explicitly introducing metapopulation structure and dynamics into the allocation problem.

LITERATURE CITED

Bodie, Zvi, Alex Kane, and Alan J. Marcus. 1996. *Investments*. 3rd edition. Chicago: Irwin.

Varian, Hal. 1993. *Microeconomic Analysis*. 3rd edition. New York: W.W. Norton.

Wu, JunJie, Richard M. Adams, and William G. Boggess. 2000. Cumulative Effects and Optimal Targeting of Conservation Efforts: Steelhead Trout Habitat Enhancements in Oregon. *American Journal of Agricultural Economics* 82(May):400-413.

LITERATURE CITED (CONT'D.)

Winston, Wayne L. 1994. *Operations Research: Applications and Algorithms*. 3rd edition. Belmont, CA: Duxbury Press.